

New simulation strategies for lattice gauge theory

Michael G. Endres*

*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139,
USA*

E-mail: endres@mit.edu

Despite the numerous successful applications of lattice QCD in nuclear and particle theory, fundamental algorithmic challenges remain. Among those, relevant for numerical studies of QCD on a space-time torus, is topological freezing—a form of critical slowing down, which becomes particularly acute for lattice spacings less than 0.05 fm. In these proceedings, I highlight several recently proposed simulation strategies for ameliorating the problem of topological freezing, and discuss both their advantages and disadvantages. Then, I turn focus toward potential strategies for addressing critical slowing down in a more general context.

*34th annual International Symposium on Lattice Field Theory
24-30 July 2016
University of Southampton, UK*

*Speaker.

1. Introduction

One of the goals of any numerical study is to provide reliable determinations of statistical and systematic uncertainties on stochastically estimated quantities. Control over both kinds of uncertainties require, in part, an understanding of correlations in the generated statistical ensembles. Such correlations arise as a consequence of the sequential nature of Markov Chain Monte Carlo (MCMC) simulations, as illustrated in Fig. 1. MCMC techniques are heavily used for numerical studies in lattice field theory, and the associated correlations are highly dependent on the time scales of the Markov process.

In general, coarse configurations generated by an MCMC algorithm will become decorrelated faster than fine configurations. The quality of an updating algorithm, however, is often determined by how well the decorrelation time for configuration generation scales with the inverse lattice spacing—a strong scaling behavior being indicative of *critical slowing down*. Observables (\mathcal{O}) serve as probes for estimating the decorrelation time of an algorithm. In particular, the integrated autocorrelation time of an observable provides a lower bound on the decorrelation time, and typically exhibits a power-law scaling behavior of the form




Figure 1: Sampling of configurations (represented by boxes, labeled by s_i for $i = 0, 1, 2, \dots$) in a configuration space (shaded region) by an MCMC algorithm.

$$\tau_{int}(\mathcal{O}) \sim \left(\frac{1}{a}\right)^{z_{int}(\mathcal{O})}, \quad (1.1)$$

where a is the lattice spacing, and $z_{int}(\mathcal{O})$ is bounded by the dynamical exponent, which is characteristic of the algorithm. Note that the integrated autocorrelation time depends on the observable in question, and in particular on how strongly the observable couples to the various modes of evolution. An explicit example of this dependence (for a variety of observables) is provided in Fig. 2 for pure $SU(3)$ gauge theory using a heat bath (HB) algorithm. The local nature of MCMC updating algorithms often (although not always; see e.g., [1]) imply $z_{int}(\mathcal{O}) \sim 2$ for local observables. This can be understood in terms of the diffusive character of the algorithm, although some quantities may have a far worse scaling behavior, as also shown in the figure.

In perhaps the most severe of forms, critical slowing down can result in the freezing of topological charge for gauge theories defined on a periodic lattice. This can be understood by noting that in the continuum, the gauge configuration space is separated into distinct topological sectors labeled by a quantized charge $Q \in \mathbb{Z}$. On the lattice, however, the configuration space is connected, albeit with large action barriers at non-integer Q . Intuitively this is due to the fact that changes in topology require introducing discontinuities in the gauge configuration, which result in a large Euclidean action. The free energy at a fixed topological charge is given by

$$f(Q) = -\frac{1}{V} \log \langle \delta(Q - Q(s)) \rangle, \quad (1.2)$$

where $\langle \dots \rangle$ is an ensemble average over decorrelated configurations s . An illustration of the free energy is provided in Fig. 3 (left); note that the height of the free-energy barriers depends on the choice of action, and diverges as the continuum limit is approached. A change in topology during an MCMC simulation requires tunneling through barriers, and the likelihood for such occurrences is therefore expected to be exponentially suppressed, particularly in the regime of fine lattice spacing. At ultra-fine lattice spacing, the time scale for such changes—controlled by the barrier height—can exceed the total feasible simulation time, and thus results in *topological freezing*. This phenomena becomes particularly acute when $a \lesssim 0.5$ fm, and can lead to systematic errors, such as power-law finite volume corrections to masses and other observables [3]¹.

Since the previous Lattice Conference, there have been a number of interesting and promising new proposals for addressing the challenge of topological freezing from the standpoint of gauge generation. These ideas fall into two classes: those that approach the problem from the infrared—by exploiting the fact that a different choice of boundary conditions can influence the connectivity of field space [11] (as motivated by [12]), and those that approach the problem from the ultra-violet—by exploiting the ambiguity in defining topological charge at finite lattice spacing [2, 13, 14]. In the next sections, I briefly review these strategies, as well as highlight some of their advantages and drawbacks (also see [15, 16, 17] for further discussion of these approaches, as presented at this conference). It should be mentioned that a number of strategies have been discussed at this conference which address topological freezing from an analysis standpoint rather from the standpoint of gauge generation [18, 19]. Although these are very interesting and import works, they will not be discussed here in any detail.

2. Non-orientable manifolds

A key insight which motivates the use of non-orientable manifolds to solve topological freezing is the observation that the topological character of field space depends not only on the gauge group but also the nature of the space-time manifold. With the use of open boundary conditions [12], for example, topological charge is no longer quantized and free-energy barriers between the different topological sectors are absent. Physically, changes in topology can occur when topo-




Figure 2: Lattice spacing dependence of the integrated autocorrelation time for pure $SU(3)$ gauge theory on a four-dimensional periodic lattice. Simulations were performed using the HB algorithm at a fixed physical volume (see [2] for details).

¹It should be noted that the volume dependence of physical observables can be described by analytic formulas (see, e.g., [4, 5] for extensions of the results obtained in [3]), and with sufficient data at multiple volumes and multiple fixed topological charge sectors, a direct determination of physical quantities at vanishing theta vacuum is possible. Demonstrations of this strategy are provided in [6]. Several other fixed topology approaches for measuring topological susceptibility exist (see, e.g., [7, 8]), and have been shown to work well under appropriate conditions [9, 10].

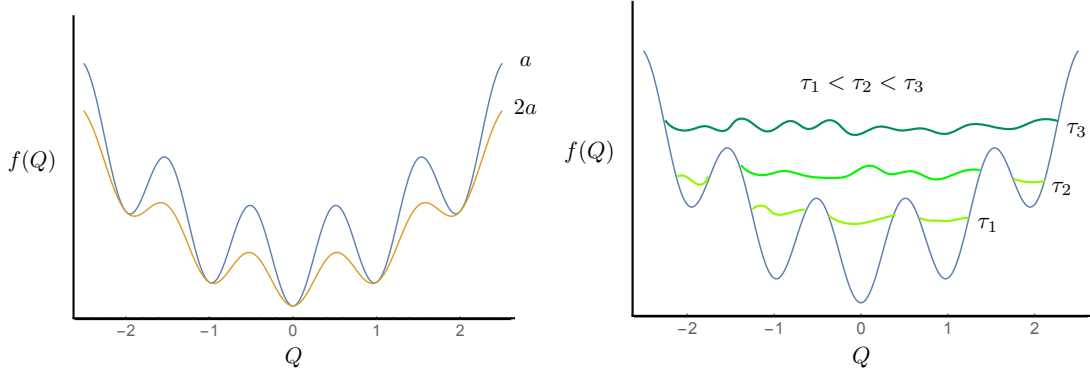


Figure 3: Left: Schematic plot of the free energy as a function of the topological charge for various lattice spacings; as the lattice spacing decreases, the height of the potential barriers diverge. Right: Schematic plot of the biased effective potential (green) for topological charge at various stages of a simulation, labeled by times τ_i ($i = 1, 2, 3$).

logical charge flows in and out of the system at the open boundaries. Although the use of open boundary conditions resolves the problem of topological freezing and have been used successfully in practical applications (see, e.g., [20]), their use comes at the expense of explicitly breaking translational symmetry. The systematic errors induced by open boundary conditions on observables (such as unphysical effects on correlation functions near the boundary), must therefore be evaluated—and if necessary—accounted for (see, e.g., [21, 22, 23, 24] for further discussion).

The use of non-orientable manifolds introduces a twist to this strategy by employing P-periodic boundary conditions (i.e., imposing a parity transformation on all fields at the boundary) as opposed to open boundary conditions in one or more of the space-time directions [11, 15]. Consequences of this choice are several-fold: 1) topological charge is no longer quantized, and 2) there is no local breaking of translational invariance and thus translational symmetry breaking effects are suppressed. Numerical studies of quenched QCD using an implementation of P-periodic boundary conditions suggest that the scaling of $\tau_{int}(Q)$ with lattice spacing is comparable to that of open boundary conditions, and thus offers an improvement compared to that of periodic boundary conditions.

Despite the promising outlook, there are a number of disadvantages to using this approach. The first is that although topology is no longer quantized, the classical field space remains broken into two distinct sectors in the continuum. Consequently, sampling of these two sectors is susceptible to freezing on the lattice at fine lattice spacing in much the same manner as the topological freezing previously discussed. The second is that although implementation of fermions on such space-time manifolds is possible, it is nonetheless somewhat nontrivial (see, e.g., [16] for further details).

3. Metadynamics

The strategy taken in metadynamics approaches is to introduce a Monte Carlo (MC) time dependent potential bias to the action, which depends on one or more collective coordinate variables. The potential bias is designed to disfavor revisiting past values of the collective coordinate during the MCMC evolution. In the case of topological charge, such a bias—call it $w(Q, \tau)$ —has the

effect of “filling” the potential wells as MC time progresses, thereby enabling eventual spill-over into different topological sectors. An illustration of this idea is provided in Fig. 3 (right) for various MC times during a hypothetical simulation. Note that at late times, the collective variable—in this case topological charge—is eventually sampled uniformly.² An important observation is that after the potential wells have been filled and the sampling of sectors becomes uniform, the time average of the potential bias provides a reliable estimator for the free energy. Specifically, one finds $f(Q) = -\bar{w}(Q)$, where $\bar{w}(Q) = N_\tau^{-1} \sum_\tau w(Q, \tau)$ and N_τ is the number of MC time samples at equilibrium. Given a reliable estimator for $f(Q)$, one may then reconstruct expectation values via reweighting

$$\langle \mathcal{O} \rangle = \frac{\langle \mathcal{O} e^{-V\bar{w}(Q)} \rangle_{\text{biased}}}{\langle e^{-V\bar{w}(Q)} \rangle_{\text{biased}}} . \quad (3.1)$$

A demonstration of the metadynamics approach to solving the exponential suppression of topology change for CP^{N-1} models was provided in [14] (see also [17]), and the generalization of this method to QCD is straightforward. There are several potential drawbacks to this strategy, however:

1. The equilibration time required to fill the potential wells scales inversely with the barrier height; this overhead can presumably be avoided with a good initial guess for the functional form of the potential bias.
2. The fraction of sampled configurations contributing meaningfully to reweighting average is expected to scale inversely with the barrier height Δf .

With regard to the latter, assuming $\Delta f \sim (1/a)^h$, the simulation cost (once equilibrated) effectively scales as

$$\text{cost} \propto \left(\frac{1}{a} \right)^{D+z_{\text{int}}(\mathcal{O})+h} , \quad (3.2)$$

where $z_{\text{int}}(\mathcal{O})$ is the dynamical exponent, determined from the autocorrelations of \mathcal{O} in the metadynamics simulation. The scaling of $z_{\text{int}}(\mathcal{O})$ was not determined in [14], but given the bias potential for Q is flat, a reasonable guess might be $z_{\text{int}}(Q) \sim 2$. Although the exponential suppression of topology change is eliminated, the approach nonetheless demands additional costs (due to the additional exponent h), which should be accounted for when comparing with other methods.

4. Multiscale thermalization

Multiscale methods have played an important role in MC simulations. For example, they have been used to improve the efficiency of Dirac operator inversion [25, 26, 27, 28], and reducing

²This assumes that a cutoff has been introduced on the magnitude of the topological charge. In practice, such a cutoff must be introduced, and if chosen sufficiently large compared to the natural width of the topological charge distribution (as governed by the chiral susceptibility and the physical spacetime volume) should induce negligible systematic artifacts.

statistical noise in estimates of correlation functions and other observables [29, 30, 31]. Implementation of a multiscale algorithm for gauge field updating in lattice QCD, however, remains an open challenge despite some early progress for some simpler field theories [32, 33, 34, 35, 36, 37].

Recently, multiscale methods have been used to rapidly thermalize gauge field configurations by combining the multigrid concepts of restriction (coarse-graining) and prolongation (refinement) with the renormalization group [2, 13]. To better understand the strategy behind this approach, it is helpful to first review some basic aspects of a MCMC simulation. A typical MC simulation begins with an initial configuration s_0 drawn from a known distribution \mathcal{P}_0 (e.g., an ordered or disordered configuration). Subsequent configurations s_τ (with $\tau > 0$) are then generated by a Markov process, defined by the transition matrix \mathcal{M} , and constructed such that the desired (target) distribution \mathcal{P} is a stationary distribution (e.g., $\mathcal{M}\mathcal{P} = \mathcal{P}$). The τ -th configuration generated in the Markov Chain will be drawn from a distribution $\mathcal{P}_\tau = \mathcal{M}^\tau \mathcal{P}_0$. Under suitable assumptions (e.g., detailed balance, ergodicity, etc.), the time evolution of the probability distribution can be expressed as

$$\mathcal{P}_\tau(s) = \mathcal{P}(s) + \sum_{n \geq 1} \langle s | \chi_n \rangle \langle \tilde{\chi}_n | \mathcal{P}_0 \rangle e^{-\tau/\tau_n}, \quad (4.1)$$

where χ_n and $\tilde{\chi}_n$ are the right and left eigenvectors of \mathcal{M} [$\chi_n(s) \equiv \langle s | \chi_n \rangle$], and τ_n are the time scales associated with the process (see, e.g., [38]). Note that the slowest time scale, τ_1 , dictates the thermalization and decorrelation time scales of the algorithm.

Multiscale thermalization addresses the following question: *can an initial distribution of configurations \mathcal{P}_0 be efficiently produced such that the overlap $\langle \tilde{\chi}_n | \mathcal{P}_0 \rangle$ vanishes for the slowest modes of evolution?* Doing so would enable more rapid thermalization, since the approach to the equilibrium distribution \mathcal{P} would be governed only by the faster times scales (e.g., $\tau_n, n > 1$). This idea is very much analogous to the common practice for extracting energies and overlap factors from correlation functions. In the context of correlators, one can construct optimized sources from a basis of operators by solving a generalized eigenvalue problem [39, 40, 41, 42]. Of those constructed, the sources with vanishing overlap onto the ground state enable reliable determinations of excited state energies [39, 40, 41, 42]. In the context of multiscale thermalization, the initial distribution is analogous to the sources, and the same orthogonality is desired. The means by which this orthogonality is realized, however, is very different.

Before discussing the construction of initial distributions \mathcal{P}_0 which are orthogonal to the slow modes of evolution, I first summarize some important general consequences of the proposed construction, under the assumption that such orthogonality is achieved:

1. The time scale for reaching equilibrium will be shorter than the full autocorrelation time scale (i.e., the longest time scale of the algorithm, τ_1).
2. The relevant equilibration time scale, although less than τ_1 , is generally unknown and depends on the quality of the projection; thus metrics for judging equilibration are necessary.
3. The ensemble obtained after thermalization *need not* be uncorrelated with the initial ensemble (the degree of correlation with the initial ensemble is irrelevant for the purpose of determining estimates, providing the resulting distribution is equilibrated). Below, I will show an explicit example where the thermalized ensemble is in fact highly correlated with the initial ensemble, yet is properly distributed according to the desired fine action.

4. The fluctuations of observables are characterized by higher moments of the observable distribution; these moments are themselves observables, and as such, will equilibrate on the same faster time scale.

Note that the degree to which each of these properties will hold depends entirely on one's ability to prepare an initial distribution of configurations which has vanishing overlap onto the slow modes of evolution.

Now that the principle idea behind the rapid thermalization method has been outlined, I will turn to the difficult task of preparing initial distributions that realize the desired properties established above. A strategy for constructing initial distributions can be motivated and illustrated by a simple toy example, the one-dimensional Ising model. The Hamiltonian for this model is given by

$$H = -J \sum_{\langle ij \rangle} s_i s_j , \quad (4.2)$$

where the sum is taken over nearest neighbor sites, and the spins associated with each site take the possible values $s_i = \pm 1$. The partition function for the theory on an finite periodic lattice is given by

$$Z = \sum_{\{s_i\}} e^{-H} = \text{Tr} W^N , \quad W = \begin{pmatrix} e^J & e^{-J} \\ e^{-J} & e^J \end{pmatrix} , \quad (4.3)$$

where N is the total number of sites. Note that one can perform a real-space RG transformation by integrating out the degrees of freedom associated with even sites, as illustrated in Fig. 4 (left). Operationally, the coarse-graining procedure amounts to expressing the partition function as

$$Z = \text{Tr} (W^2)^{N/2} , \quad W^2 \propto \begin{pmatrix} e^{R(J)} & e^{-R(J)} \\ e^{-R(J)} & e^{R(J)} \end{pmatrix} . \quad (4.4)$$

where N is taken to be even and

$$R(J) = \frac{1}{2} \log \cosh(2J) . \quad (4.5)$$

Note that for this particular system, a comparison of Eq. 4.3 and Eq. 4.4 reveals that the coarse-grained Hamiltonian is exactly identical in form to the original Hamiltonian (up to an irrelevant constant), however, the coupling constant undergoes the transformation $J \rightarrow R(J)$. Performing such a transformation n times (assuming N is divisible by 2^n) yields a coarse-grained Hamiltonian of the form

$$H_n = -J_n \sum_{\langle ij \rangle} s_i s_j , \quad (4.6)$$

where $J_{n+1} = R(J_n)$ and $J_0 = J$.

Next imagine an ensemble of coarse configurations which have been generated according to the Boltzmann weight defined by a coarse Hamiltonian. One may ask whether it is possible to “integrate in” the remaining degrees of freedom, thereby obtaining a fine ensemble properly distributed according to the corresponding fine Hamiltonian. For this example, the procedure is




Figure 4: Left: Every other degree of freedom is integrated out, leaving behind a theory defined on a coarse lattice. Right: Coarse degrees of freedom are mapped to a fine lattice, and the remaining fine degrees of freedom are determined stochastically, using the probability measure $e^{-\beta H_n}$.

straight-forward. For every configuration within the coarse ensemble, one may first map the coarse degrees of freedom onto the fine lattice, as illustrated in Fig. 4 (right). Subsequently, the remaining degrees of freedom may be determined using a single HB update per site, using the fine Hamiltonian defined at the given refinement level. This strategy may be iterated repeatedly over multiple levels of refinement along the renormalization group (RG) flow, starting from a solitary site, as illustrated in Fig. 5. Doing so yields fully decorrelated and completely thermalized ensembles at the finest level, and the computational cost is only a single HB update per site per configuration.

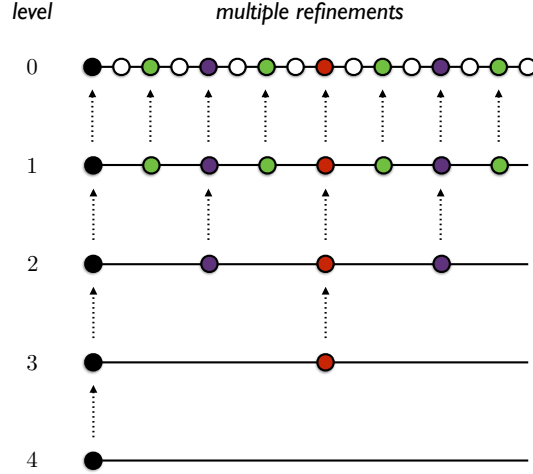


Figure 5: Generation of a lattice configuration via multiple refinement steps. Configurations at a coarse level $n + 1$ are mapped to a lattice at level n , and the remaining degrees of freedom are subsequently determined stochastically. The process is repeated until reaching the top level.

Generalization of this strategy to higher dimensions, and to more complicated systems is highly nontrivial. In all but the simplest cases, the RG matched coarse action will involve increasingly complicated and nonlocal interactions at each subsequent stage of coarse-graining. Further complicating the matter, fermion interactions are usually represented by nonlocal fermion determinants. Finally, an exact refinement prescription—as outlined in the toy model above—may no longer be feasible due to the coupling of interpolated degrees of freedom at the refined level (the one-dimensional Ising model is a special case, where interpolated degrees of freedom only couple to coarse degrees of freedom, which have been mapped to the fine lattice). One way to proceed in these cases is to rely upon a variety of approximations. These include:

1. Truncating the coarse action to only its dominant interactions; an example of such a truncation for the case of a two-dimensional Ising model is shown in Fig. 6.
2. Using an inexact refinement prescription; for example, for systems involving continuous degrees of freedom, one might consider a refinement prescription based on interpolation.

These approximations, by construction, primarily manifest as short-distance errors in the resulting fine ensemble. These errors may be corrected by *rethermalizing* the ensemble using conventional MC algorithms. Although conventional algorithms based on local updating have rather long decorrelation times, one might expect (and one can explicitly verify) that the rethermalization time is rapid by comparison since it is only the short distance properties of the ensemble that require correction. Connecting back to Eq. 4.1, one sees that generation and refinement of RG matched coarse ensembles provides a means by which to realize an initial distribution of fine configurations \mathcal{P}_0 , and the rapid nature of the rethermalization implies that this distribution is nearly orthogonal to the slow modes of evolution, as desired.




Figure 6: Left: A schematic illustration of the interactions (nearest, next-to-nearest and plaquette) induced by coarse-graining a two-dimensional Ising model (open circles represent degrees of freedom that have been integrated out). Right: Interactions retained after truncating the coarse-grained Hamiltonian.

To summarize the discussion thus far, based on the insights of the one-dimensional Ising model, the proposed multiscale thermalization scheme entails the following set of steps:

1. Generate a coarse ensemble using a RG matched coarse action.
2. Map the coarse ensemble onto a fine lattice, *while preserving the long distance properties of the ensemble*.
3. Rethermalize and evolve multiple refined streams using the fine action to correct any errors induced due to approximations.

In Fig. 7, I provide a schematic comparison of conventional and proposed approaches. In the conventional approach, the computational cost per configuration is controlled by the decorrelation time $\tau_{decorr} \sim \tau_1$. In the multiscale thermalization approach, the computational cost is controlled by the rethermalization time $\tau_{retherm}$; ideally, the coarse matching should be chosen and refinement prescription should be designed such that $\tau_{retherm} \sim \tau_k < \tau_1$ for some mode of evolution satisfying $k > 1$.

I now turn toward a specific realization of the above proposal for gauge theories. There are a number strategies for mapping coarse configurations onto a fine lattice. The most natural strategy for continuous degrees of freedom is to simply interpolate the fields, as previously mentioned (also




Figure 7: A schematic illustration of a conventional MC simulation (a) and the strategy taken for multiscale thermalization (b). Boxes represent coarse (\square) and fine (\boxplus) configurations which are either unthermalized (empty) or thermalized (gray filled). Dashed lines enclose the thermalized ensemble used to perform estimates. Note that in a conventional simulation, the number of MC updates required to obtain a decorrelated configuration in equilibrium is τ_{decorr} , whereas for multiscale thermalization it is τ_{retherm} .

see, e.g., [43, 44, 45]); a method inspired by 't Hooft [45] was applied successfully for this purpose in [2]. The benefits of the 't Hooft-inspired scheme are that the implementation is simple, local, and efficient. Furthermore, it preserves the hypercubic and translational symmetries of the coarse lattice. Finally, and most importantly, the scheme preserves the long-distance properties of the coarse configurations. In particular, the scheme preserves all even length Wilson loops originating from even sites, as well as the topological charge, and topological charge density at sufficiently fine lattice spacing. In Fig. 8, an explicit example of this preservation property is shown, where a fine $SU(3)$ gauge configuration had been coarse-grained and refined, and the topological charge density subsequently compared. One can see that indeed the gross structure is identical between fine and refined configurations, and it is only the short distance structure that requires correction. In Fig. 9 (left), a plot of the correlation between fine and refined topological charge is displayed as a function of the lattice spacing for pure $SU(3)$ gauge theory. Although at coarse lattice spacing the correlation is quite poor, it rapidly rises towards unity for lattice spacings below 0.7 fm.

In [2] and [13] the quality of the initial distributions obtained by prolongation of matched coarse ensembles was studied extensively in terms of the distribution overlap onto the slow modes of evolution. Long distance observables serve as probes for determining this overlap, as made evident by the time scale required to equilibrate the ensemble by conventional fine evolution. These time scales were studied using a variety of observables, including the action density evaluate on fields at large Wilson flow times [46], and the pion correlation function estimated at half the temporal extent of the lattice (T). Studies were performed for both pure $SU(3)$ gauge theory, as demonstrated in Fig. 9 (right), and for two color QCD with two heavy quark flavors [$am_\pi = 0.4446(6)$ and $af_\pi = 0.1001(18)$], as demonstrated in Fig. 10. In both cases the thermalization time required




Figure 8: Top: Coarse-graining and refinement of the degrees of freedom associated with the bonds of a space-time lattice. Bottom: Topological charge density associated with a smooth lattice field configuration (left) and that obtained from the same configuration after coarse-graining and then subsequent refinement (right).

for simulations starting from hot (disordered) and cold (ordered) configurations, and the associated rethermalization time required for configurations prepared by prolongating RG matched coarse ensembles are displayed. In both the pure gauge theory, and the theory with dynamical fermions, the rethermalization times appear to be rapid compared to typical thermalization times for hot and cold starts. Furthermore, the rethermalization time for non-topological quantities appear to be far shorter than the decorrelation times for a conventional simulation, as probed by observables sensitive to the topological charge. In both cases, the algorithm employed for (re)thermalization was a conventional Hybrid Monte Carlo (HMC) algorithm. Note that within this framework for constructing initial distributions \mathcal{P}_0 , the overlap with the slowest mode of evolution will not vanish exactly due to the fact that the topological charge distribution of the initial ensemble has lattice artifacts inherited from the coarse action. However, these lattice artifacts may be controlled and systematically removed with the use of improved actions at the coarse level (or possibly by subsequent reweighting).

In the case of dynamical QCD, understanding the properties of the Dirac spectrum on refined ensembles played a crucial role in the successful application of the multiscale thermalization method [13]. To understand why, it is helpful to compare the Hermitian Dirac spectrum for both the RG matched coarse ensemble, shown in Fig. 11 (left), and the spectrum computed on the corresponding fine ensemble obtained using the same (prolongated) coarse configurations, shown in Fig. 11 (right, labeled $R^{(0)}$). In the latter case, one finds a significant accumulation of spurious near-zero modes, which can be attributed to the short-distance character of the refined ensemble. These near-zero modes are particularly problematic for the rethermalization stage of the multiscale approach, since the large fermion forces they induce can lead to instabilities in the evolution. However, it was found that a gap in the Dirac spectrum could be produced by performing a very short quenched evolution of the refined ensemble (e.g., two HMC trajectories), after which, proper dynamical gauge evolution could proceed without issues. In Fig. 11 (right, labeled $R^{(1)}$ and $R^{(2)}$)

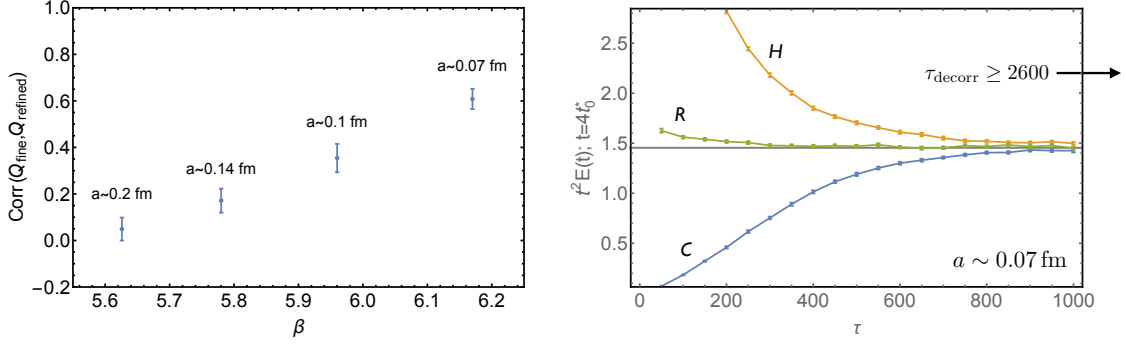


Figure 9: Left: Correlation between the topological charge for configurations before and after coarse graining and refinement (as depicted in Fig. 8) as a function of the lattice spacing for pure $SU(3)$ gauge theory defined by the Wilson gauge action. Right: Plot of thermalization curves for hot (H) and cold (C) starts, and for an ensemble generated using a RG matched coarse action and subsequently refined (R).




Figure 10: Rethermalization curves for Yang-Mills action density (left) and pion correlator at large time separation (right) as a function of the number of HMC trajectories for $SU(2)$ gauge theory with two heavy quark flavors. Curves represent thermalization from a cold start (C), a hot start (H), and refined ensemble after evolving for two quenched HMC trajectories ($R^{(2)}$).

examples of Dirac spectra after one and two quenched trajectories are shown, respectively.

5. Conclusion

Several interesting and promising strategies for addressing topological freezing have been proposed during the past year. In this talk, I have introduced the basic concepts behind three ideas that tackle the problem from the standpoint of ensemble generation, namely the use of non-orientable manifolds, metadynamics simulations, and multiscale thermalization. It should be emphasized that each of these methods provide very new and very different approaches toward addressing the same underlying problem. Although each of these strategies show significant promise based on exploratory studies, they also possess some rather interesting drawbacks and opportunities for further exploration and improvement.

Although the results for multiscale thermalization appear very promising, the success of that strategy relies upon the ability to prepare an initial distribution (\mathcal{P}_0) of configurations that is or-




Figure 11: Left: Hermitian Dirac spectrum for $SU(2)$ gauge theory with two heavy quark flavors, calculated on a RG matched coarse ensemble. Right: Hermitian Dirac spectrum for the same theory, determined on the corresponding fine lattice using refined field configurations, $R^{(0)}$. Also shown is the Dirac spectrum after several quenched HMC steps, $R^{(\tau)}$, for $\tau > 0$. The vertical dashed line corresponds to the bare PCAC mass determined on an equilibrated fine ensemble.

thogonal to the lowest modes of evolution. As had been demonstrated, one successful realization of this strategy is based upon refinement of a coarse ensemble that has been generated using a RG matched coarse action. The advantages and disadvantages of that realization in the context of gauge theories are detailed at length in [2]. These include:

1. The initial ensemble obtained by prolongation will have a topological charge distribution that is only correct up to coarse action lattice artifacts (removal of those lattice artifacts require rethermalization time scales that are multiples of a full autocorrelation time).
2. The coarse lattice artifacts inherited by the prolonged ensemble, however, can be systematically removed by improving the coarse lattice action (as is often the case, the tuning involved requires some additional overhead cost).
3. Non-topological quantities (measured at fixed topology, or on a fixed but possibly incorrect topological charge distribution) appear to rethermalize on time scales that are significantly shorter than the full autocorrelation time, based on empirical evidence (i.e., by studying rethermalization times of various long distance observables).

As discussed earlier, if an initial configuration distribution has vanishing overlap with the slow modes of evolution, then the equilibration time will be controlled by shorter time scales. As such, the full autocorrelation time—as governed by the slowest mode of evolution—is no longer a relevant time scale in the problem. This observation raises an interesting question: *can multiscale thermalization provide a new avenue for addressing the problem of critical slowing down more generally (i.e., beyond topological freezing), by avoiding the problem altogether?* To clarify the question, consider setting aside issues of topology—either by choice of boundary conditions (open or P-periodic) or use of a different algorithm (e.g., metadynamics). Intuition suggests that interpolated coarse RG matched ensembles will yield better approximations to the fine thermalized ensemble as the continuum limit is approached. This follows from the simple observation that one can more reliably interpolate a smooth field compared to a rough field. It is possible that an increased number

of modes of evolution will be projected out as a result of this procedure as one approaches the continuum limit, since fields in that regime become locally flat. The interplay between the number of modes projected out and the scaling of excited modes with lattice spacing as one makes this approach is an interesting and open question. It may very well be that the rethermalization time required in the multiscale approach scales more slowly than the decorrelation time for conventional algorithms as a function of lattice spacing. If this is indeed the case, then multiscale thermalization would likely provide a new path for avoiding critical slowing down more generally.

Acknowledgments

M.G.E. would like to thank R. C. Brower, W. Detmold, K. Orginos and A. V. Pochinsky for informative discussions, and W. Bietenholz for several constructive comments in response to this presentation. This work was partially supported by the U. S. Department of Energy through Early Career Research Award de-sc0010495 and under Grant No. de-sc0011090.

References

- [1] N. Prokof'ev and B. Svistunov, *Phys. Rev. Lett.* **87**, 160601 (2001).
- [2] M. G. Endres, R. C. Brower, W. Detmold, K. Orginos, and A. V. Pochinsky, *Phys. Rev.* **D92**, 114516 (2015), 1510.04675.
- [3] R. Brower, S. Chandrasekharan, J. W. Negele, and U. J. Wiese, *Phys. Lett.* **B560**, 64 (2003), hep-lat/0302005.
- [4] A. Dromard, W. Bietenholz, U. Gerber, H. Mejía-Díaz, and M. Wagner, *Acta Phys. Polon. Supp.* **8**, 391 (2015), 1505.03435.
- [5] A. Dromard, W. Bietenholz, U. Gerber, H. Mejía-Díaz, and M. Wagner, *PoS LAT2015*, 278 (2016), 1510.08809.
- [6] W. Bietenholz *et al.*, *Phys. Rev.* **D93**, 114516 (2016), 1603.05630.
- [7] S. Aoki, H. Fukaya, S. Hashimoto, and T. Onogi, *Phys. Rev.* **D76**, 054508 (2007), 0707.0396.
- [8] P. de Forcrand *et al.*, *Nucl. Phys. Proc. Suppl.* **73**, 578 (1999), hep-lat/9810033.
- [9] W. Bietenholz, P. de Forcrand, and U. Gerber, *JHEP* **12**, 070 (2015), 1509.06433.
- [10] I. Bautista *et al.*, *Phys. Rev.* **D92**, 114510 (2015), 1503.06853.
- [11] S. Mages *et al.*, (2015), 1512.06804.
- [12] M. Lüscher and S. Schaefer, *JHEP* **07**, 036 (2011), 1105.4749.
- [13] W. Detmold and M. G. Endres, (2016), 1605.09650.
- [14] A. Laio, G. Martinelli, and F. Sanfilippo, *JHEP* **07**, 089 (2016), 1508.07270.
- [15] S. Mages, *PoS LAT2016*, 338 (2016).
- [16] B. Toth, *PoS LAT2016*, 333 (2016).
- [17] F. Sanfilippo, *PoS LAT2016*, 274 (2016).
- [18] W. Bietenholz, K. Cichy, P. de Forcrand, A. Dromard, and U. Gerber, *PoS LAT2016*, 321 (2016).

- [19] C. Bernard and D. Toussaint, PoS **LAT2016**, 189 (2016).
- [20] M. Cè, M. F. Garcia Vera, L. Giusti, and S. Schaefer, PoS **LAT2016**, 350 (2016).
- [21] M. Luscher and S. Schaefer, Comput. Phys. Commun. **184**, 519 (2013), 1206.2809.
- [22] M. Bruno, P. Korcyl, T. Korzec, S. Lottini, and S. Schaefer, PoS **LAT2014**, 089 (2014), 1411.5207.
- [23] A. Amato, G. Bali, and B. Lucini, PoS **LAT2015**, 292 (2016), 1512.00806.
- [24] B. Lucini, C. McNeile, and A. Rago, PoS **LAT2015**, 102 (2016), 1511.09303.
- [25] R. Babich *et al.*, Phys. Rev. Lett. **105**, 201602 (2010), 1005.3043.
- [26] R. Babich *et al.*, PoS **LAT2009**, 031 (2009), 0912.2186.
- [27] A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann, SIAM J. Sci. Comput. **36**, A1581 (2014), 1303.1377.
- [28] J. Brannick *et al.*, Numer. Math. (2015), 1410.7170.
- [29] M. Lüscher and P. Weisz, JHEP **09**, 010 (2001), hep-lat/0108014.
- [30] M. Cè, L. Giusti, and S. Schaefer, (2016), 1601.04587.
- [31] M. G. Vera and S. Schaefer, (2016), 1601.07155.
- [32] J. Goodman and A. D. Sokal, Phys. Rev. Lett. **56**, 1015 (1986).
- [33] R. G. Edwards, J. Goodman, and A. D. Sokal, Nucl. Phys. **B354**, 289 (1991).
- [34] R. G. Edwards, S. J. Ferreira, J. Goodman, and A. D. Sokal, Nucl. Phys. **B380**, 621 (1992), hep-lat/9112002.
- [35] W. Janke and T. Sauer, Nucl. Phys. Proc. Suppl. **34**, 771 (1994), hep-lat/9312043.
- [36] M. Grabenstein and B. Mikeska, Nucl. Phys. Proc. Suppl. **34**, 765 (1994), hep-lat/9311021.
- [37] M. Grabenstein and K. Pinn, Phys. Rev. **D50**, 6998 (1994), hep-lat/9406013.
- [38] M. Lüscher, Computational Strategies in Lattice QCD, in *Modern perspectives in lattice QCD: Quantum field theory and high performance computing. Proceedings, International School, 93rd Session, Les Houches, France, August 3-28, 2009*, pp. 331–399, 2010, 1002.4232.
- [39] C. Michael and I. Teasdale, Nucl. Phys. **B215**, 433 (1983).
- [40] C. Michael, Nucl. Phys. **B259**, 58 (1985).
- [41] M. Luscher and U. Wolff, Nucl. Phys. **B339**, 222 (1990).
- [42] B. Blossier, M. Della Morte, G. von Hippel, T. Mendes, and R. Sommer, JHEP **04**, 094 (2009), 0902.1265.
- [43] M. Lüscher, Commun. Math. Phys. **85**, 39 (1982).
- [44] A. Phillips and D. Stone, Commun. Math. Phys. **103**, 599 (1986).
- [45] G. 't Hooft, Physics Letters B **349**, 491 (1995).
- [46] M. Lüscher, JHEP **08**, 071 (2010), 1006.4518, [Erratum: JHEP03,092(2014)].